ARMY RESEARCH LABORATORY

Remote Scientific Visualization Using the Internet Protocol

by John M. Vines

ARL-TR-3609 **September 2005**

NOTICES

Disclaimers

Army Research Laboratory
Aberdeen Proving Ground, MD 21005-5067

ARL-TR-3609 **September 2005**

Remote Scientific Visualization Using the Internet Protocol

John M. Vines
Computational and Information Sciences Directorate, ARL

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE *(DD-MM-YYYY)* September 2005	**2. REPORT TYPE** Final	**3. DATES COVERED (From - To)** 5 May 2004–5 June 2004

4. TITLE AND SUBTITLE Remote Scientific Visualization Using the Internet Protocol	**5a. CONTRACT NUMBER**
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S) John M. Vines	**5d. PROJECT NUMBER** 5UH7FL
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Research Laboratory AMSRD-ARL-CI-HC Aberdeen Proving Ground, MD 21005-5067	**8. PERFORMING ORGANIZATION REPORT NUMBER** ARL-TR-3609
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	**10. SPONSOR/MONITOR'S ACRONYM(S)**
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Historically, when a user is required to work on a remote UNIX workstation, the user will log into the remote system and via the XServer and X11 protocol, set the display back to the local host. Problems arise with the requirement of high-bandwidth network requirements and the lack of OpenGL performance. Newly developed technologies such as the Teraburst V2D hardware transmitters and receivers leverage state-of-the-art compression algorithms, providing users with the performance of local graphics from a remote system.

15. SUBJECT TERMS

VideoOverIP, MPEG compression, remote display, Teraburst V2D

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON** John M. Vines
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			**19b. TELEPHONE NUMBER** (*Include area code*)
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL	20	(410) 278-9150

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures

List of Tables

INTENTIONALLY LEFT BLANK.

Acknowledgments

The author would like to thank the U.S. Army Research Laboratory Major Shared Resource Center government and contract staff and Teraburst Networks, Incorporated (now IP Video Systems Inc.) for providing technical details and information documented in this report.

1. Introduction

Until recently, the only way for researchers and engineers to see the display from a remote computational resource running an X Server was to set the display back to the local X Window client. Performance was an important aspect in the initial design of the X Window System[1] protocol and continues to be important in applications and extension development. However, sending a large computer graphics video display across a network connection greatly reduces performance and requires much more network bandwidth.

The advent of new compression algorithms and techniques as well as the development of Internet protocol-based communication technologies is greatly increasing the capacity to communicate over geographically dispersed areas in real time and ultimately changing the way we do business.

2. State-of-the-Art Technology

The U.S. Army Research Laboratory (ARL) uses VideoOverIP technology to provide real-time internet protocol (IP)-based communications between two reality centers—one based at Aberdeen Proving Ground, MD, and the other at ARL headquarters in Adelphi, MD. Personnel are able to communicate between sites by simply walking into the room, looking at a camera, and talking.

The systems are more prevalently used for collaborative efforts between personnel located at ARL and remote scientists and researchers. The input to the VideoOverIP system is generated by one of ARL's scientific visualization computational assets. The real-time graphics signal is generated at high-definition video resolution and then down-converted to NTSC via real-time down-conversion hardware (figure 1).

A remote user enters an IP address to a PC client application to view the Motion Picture Experts Group (MPEG) video and audio stream. The stream may be sent to the user via multicast, for multiple collaborative sessions at once, or unicast to enable a one-to-one session.

During an interactive collaborative session (figure 2), ARL personnel are able to make changes to scientific visualizations on the fly, emphasizing areas of interest, tweak visualization algorithms, or make changes to color tables. This can all be done while remote researchers and engineers are watching.

[1]X Window System Network Performance. http://keithp.com/~keithp/talks/usenix2003/html/net.html (accessed May 2004).

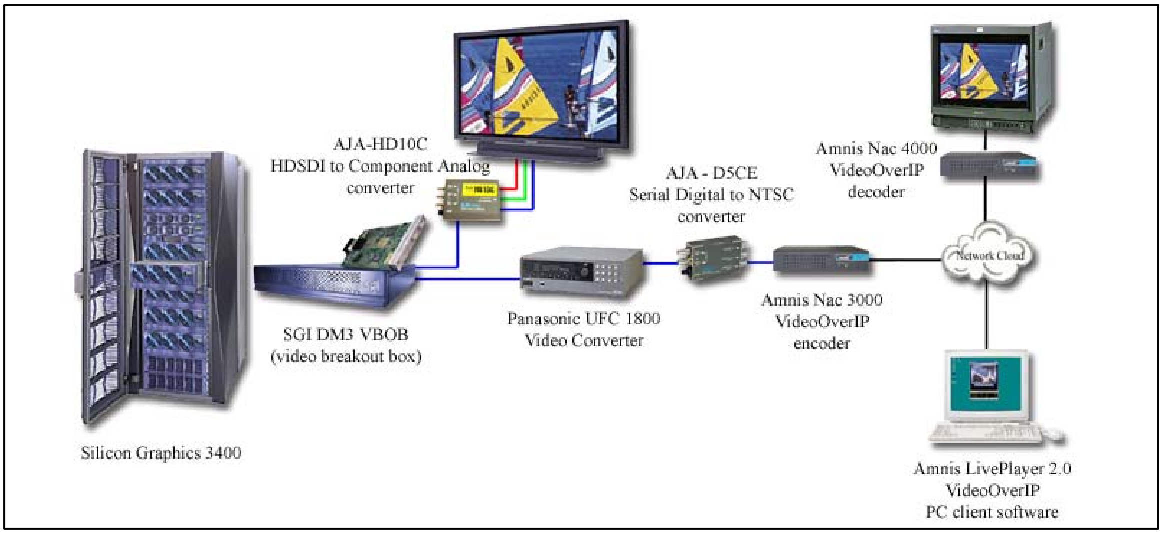

Figure 1. High-definition graphics output to standard definition VideoOverIP client.

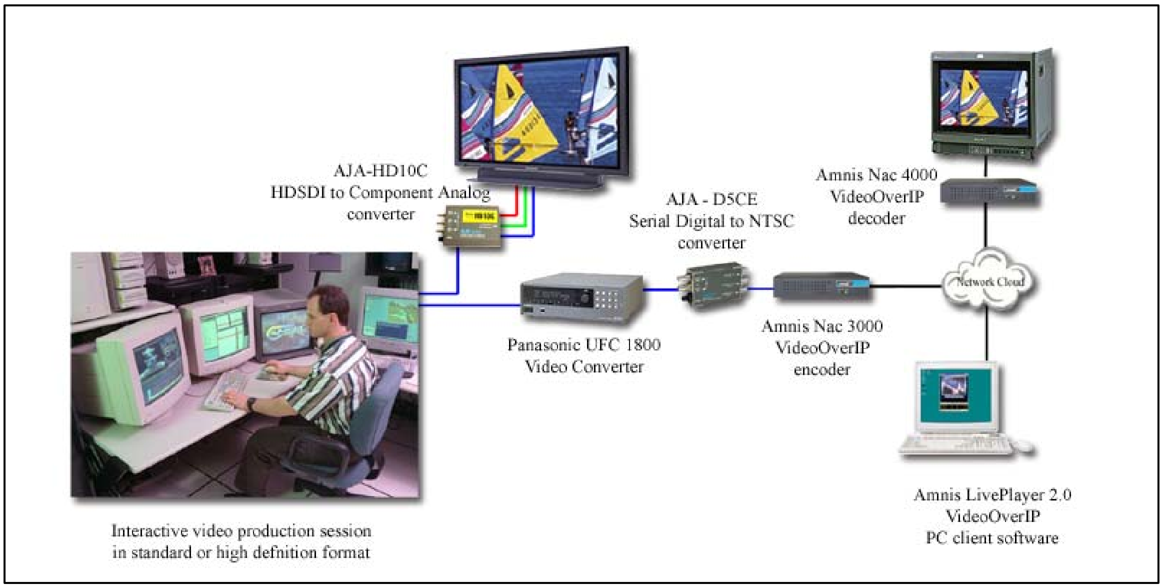

Figure 2. Interactive video production session.

The ARL Major Shared Resource Center (MSRC) has also incorporated VideoOverIP technology in the video production suite. Production personnel have the ability to establish interactive editing sessions with remote scientists and researchers to create high-definition or standard-definition videos for technical reports, briefings, and conferences.

Bandwidth requirements for VideoOverIP sessions are a function of the client network connection and client infrastructure. A VideoOverIP set-top box decoder provides the capability to decode multiple megabit MPEG 1/2/4 streams with low latency in real time. In nearly all sessions, the set-top decoders are set to decode the highest quality possible, MPEG2/4 full D1 resolution, 720 × 480, with accompanying stereo audio.

Many PC software clients require a Pentium III-class CPU running at 600 MHz or better for MPEG2 decoding. The default MPEG2 session would incorporate a MPEG stream with a resolution of 640 × 480 at 4 Mbits/s. The PC clients that cannot support the MPEG2 requirements will have the ability to decode an MPEG1 stream. ARL has successfully decoded an MPEG1 stream with a Pentium-class CPU running at 100 MHz.

MPEG1 provides resolution of ~352 × 240, 30 frames/s at 1.5 Mbits/s, and a picture of VHS quality with accompanying stereo audio.

The MPEG2 format can provide better resolution and higher quality at higher bit rates. In most instances, ARL MSRC collaborative sessions encode a MPEG2 at a resolution of 640 × 480, 30 frames/s, requiring 4 Mbits client-side bandwidth and processing power. The quality of the video is near DVD quality, with stereo audio encoded at 44.1 or 48 KHz.

VideoOverIP, simply referred to as HDOverIP, also operates at high-definition resolutions. The equipment configuration for the HDOverIP system requires the use of high-definition video production, encoding, and decoding equipment. However, the process parallels that of the standard definition VideoOverIP process.

To produce a high-definition signal, as in the case at ARL, we use one of our scientific visualization assets with an attached SGI video breakout box. We then send this high-definition signal to an encoder, which, in turn, sends the signal to a HDOverIP server and the stream to the defined IP address (figure 3). The client-side requirements to decode the high-definition signal are 19–24 Mbits of bandwidth. The host processor varies, depending upon hardware-assisted decoding or purely software decoding. A Pentium III-class at 1.0 GHz or better is a good place to start.

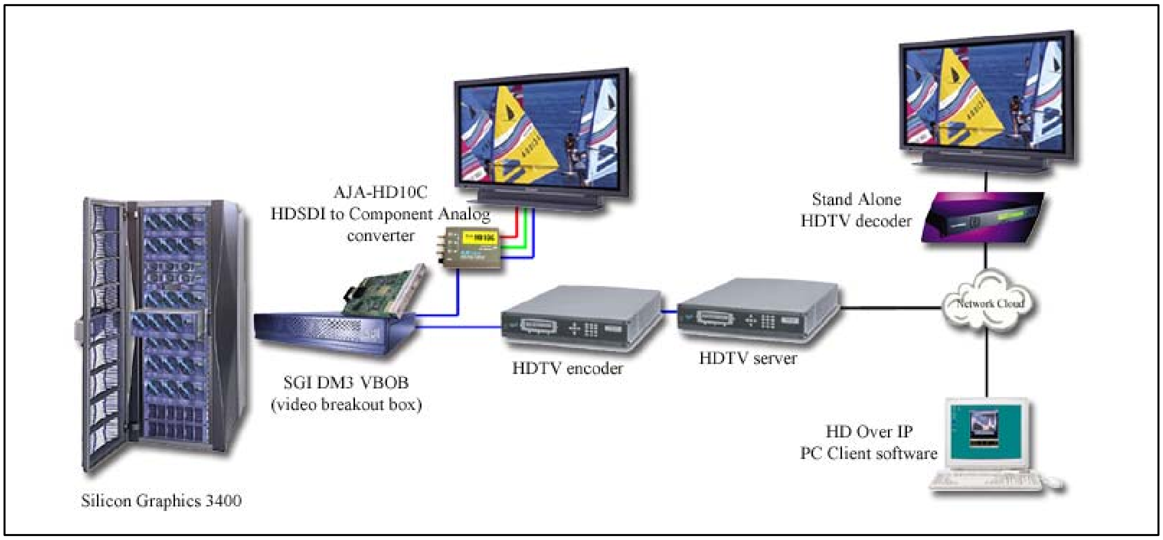

Figure 3. High-definition graphics output to high-definition HDOverIP client.

The HDOverIP technology lends itself to a computer-graphics environment. The standard supports progressively scanned video signals and the 720p, 1280 pixels wide × 720 pixels high, is very close to SXGA (1280 pixels wide × 1024 pixels high) screen resolution, therefore providing a much higher resolution image.

Formats for high-definition video include 720p (1280 pixels wide × 720 pixels high) and 1080i/p (1920 pixels wide × 1080 pixels high). A 720p signal generates 720 progressively scanned lines; 1080i generates 1080 interlaced lines; and 1080p provides 1080 progressively scanned lines. NTSC quality video produces 525 interlaced lines at a resolution of 640 pixels wide × 480 pixels high.

Using high-definition video signals and HDOverIP communication equipment lessens the artifacts and noise introduced to the video signal through the scan-converting process (figure 4). The inherent problems of scan converting a video signal are much less problematic because less pixels are squeezed into a small space of 640 × 480 pixels.

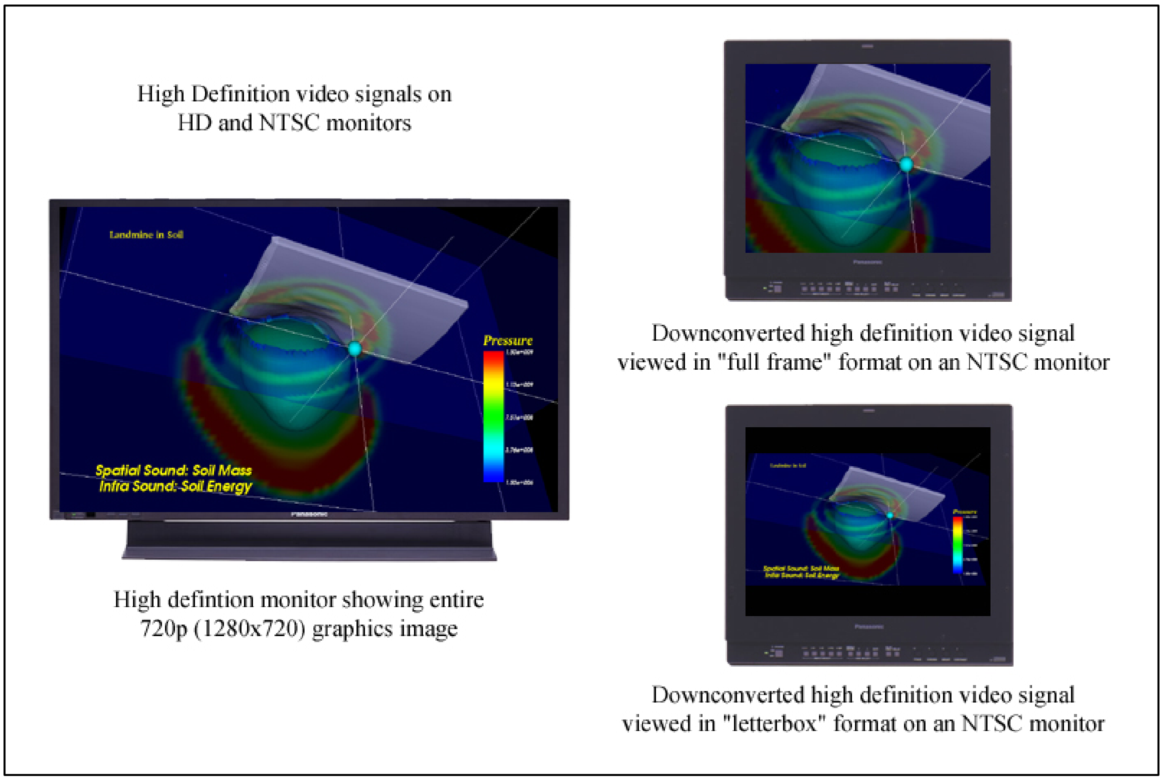

Figure 4. High-definition video signals on HD and NTSC monitors.

Today's datasets are growing to vast proportions. Gigabytes of data are quite common and growing more toward terabytes, and in the not-so-distant future, petabytes. The ability to visualize this amount of data is not a trivial task, and, in most cases, requires high-performance computing (HPC) resources with a huge amount of memory and disk space. Obviously, machines of this nature are not a common commodity and are housed at strategic sites around the world.

In most instances, the scientists and researchers running the codes that generate these vast amounts of data are not local to the computational assets and rely on networked graphic connections to view the data. The performance of the networked graphics is a function of the network. Bigger, faster, and better networks provide better performance, while smaller and slower network connections provide less performance. Unfortunately, in some cases, a network may not have the bandwidth to support the networked computer graphics and require a researcher to travel to the remote site to visualize the data.

The VideoOverIP and HDOverIP technologies enable collaboration across IP-based networks at relatively low bandwidths. In most instances, these technologies can be configured to support remote researchers, even at the lowest of bandwidths. An Ethernet connection, switched at 10baseT speeds, provides enough bandwidth for MPEG2 video at standard definition resolution. A fast-Ethernet connection provides enough bandwidth for MPEG2 video at high-definition resolutions. But neither technology addresses computer graphics video in its native resolution.

The next step in IP communication technology must be able to handle computer graphics video in its native resolution, whether it is VGA (640 × 480), SVGA (800 × 600), XGA (1024 × 768), or SXGA (1280 × 1024), and have the ability to evolve to UXGA (1600 × 1200), UXGA-W (1920 × 1200), and ultimately, QXGA (2048 × 1536).

The major applications of such a technology are sharing HPC resources, broadcast/multicast graphics, and video; connecting immersive three-dimensional (3-D) visualization centers and graphics displays and sources; and collaborating desktops between networked users.[2]

3. Cutting Edge Technology

Teraburst Networks, Inc. (now IP Video Systems Inc.) has been at the forefront of optically transmitted computer graphics video, with accompanying keyboard and mouse, for a number of years. The company has now developed a video-to-data (V2D) product which transmits over fast (100 Mb/s) or gigabit (1000 Mb/s) Ethernet/IP networks. At present, supported resolutions are UXGA (1600 × 1200) at 60 Hz and SXGA (1280 × 1024) at 96 Hz. Stereoscopic video support is provided at the SXGA resolution.

The Teraburst V2D equipment incorporates the advantages of optical networks and solutions with the flexibility of Ethernet and data networks. The devices support LAN and WAN distances, provide fully interactive synchronized video and control signals, support IP multicast and can provide further network security and data encryption by using IP VPN.

[2] Teraburst Networks, Inc. *Visualization With High-Resolution and Low-Latency Over Packet Networks*; White Paper; Sunnyvale, CA, August 2003.

A significant capability within itself is the use of an Ethernet/IP network. Packet-based networks are more available than optical-based networks; therefore, the V2D devices are better suited for broader deployment as well as making it easier to integrate desktop software clients (figures 5 and 6).

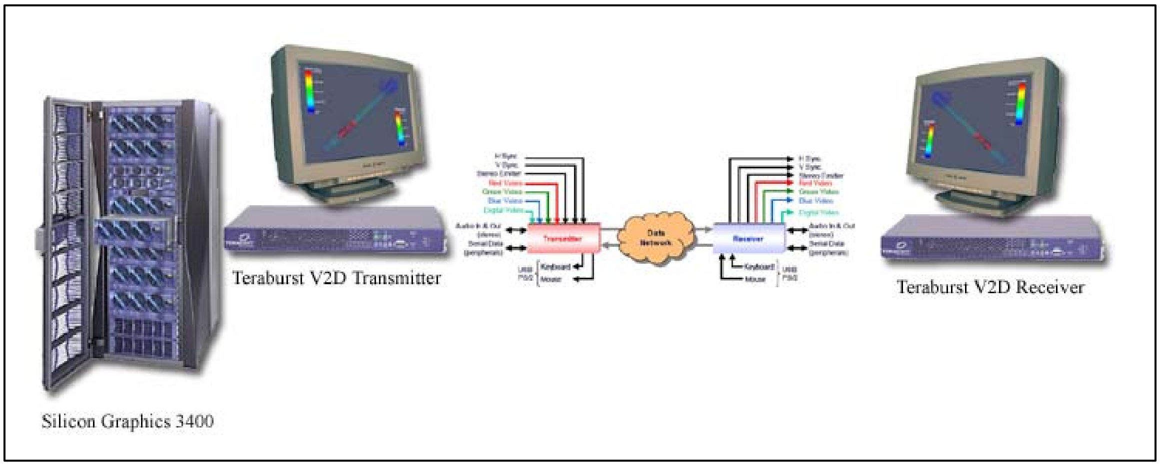

Figure 5. Teraburst Networks, Inc.'s V2D networked video transmitter and client-side receiver.

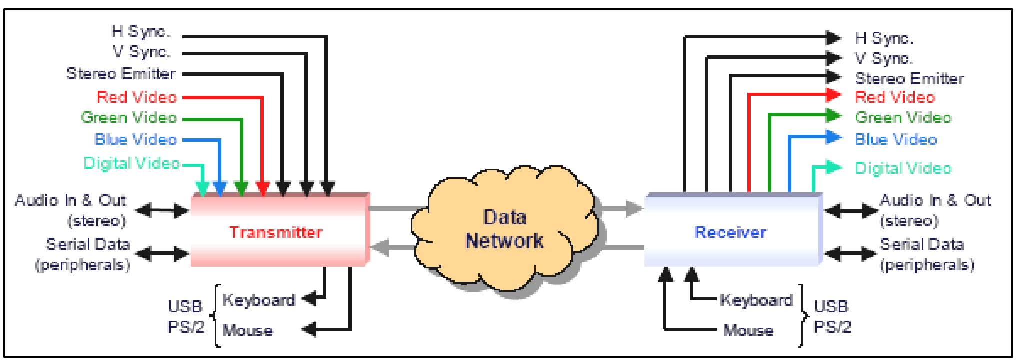

Source: Teraburst Networks, Inc.

Figure 6. Communication between Teraburst V2D equipment.

From an economical standpoint, the data products are more cost effective, simply because they do not rely on an optical network. Furthermore, IT personnel are more familiar with packet than optical technologies and have more tools available to manage and diagnose any problems.

Packet-based solutions cost less and are easier to deploy but have to deal with nondeterministic bandwidth availability and require more traffic engineering of links over which the traffic will be passing. When packets are dropped, the infrastructure must be able to compensate with mechanism based on frame and subframe synchronization to keep the video running smoothly with a minimum of artifacts. Teraburst algorithms are designed to cope with very high packet loss without losing connectivity, something which is not possible with MPEG technology.

Different compression parameters can be configured for moving images and still images, while detection of these images is fully automatic. This feature can be used to optimize required

bandwidth and needed frames per second during moving images while preserving very high quality for still images.

Compression adapted for visualization was a requirement for data networks. The main goal of compression is to reduce the required bandwidth without sacrificing overall picture quality. At present, stereo visualization applications are running at 20–30 frames per second at 96 Hz and moving toward 50–60 frames per second at 120 Hz. The screen pixel content will increase from 1 megapixel to more than 3 megapixels.[2] At 24 bits per pixel, this means transmission rates of uncompressed images will increase from 2 Gb/s to nearly 10 Gb/s.[1] MPEG2 provides highly compressed video but does not fulfill quality requirements for high-end visualization.

Teraburst has developed a compression algorithm that incorporates frame and slice differencing with variable noise reduction as the main mechanism to eliminate redundant information between frames (figures 7 and 8). This technology, referred to as Teraburst Intelligent Frame Differencing,* takes into account noise introduced by components, analog-to-digital conversion, system electronics, source stability, and switching and cabling.

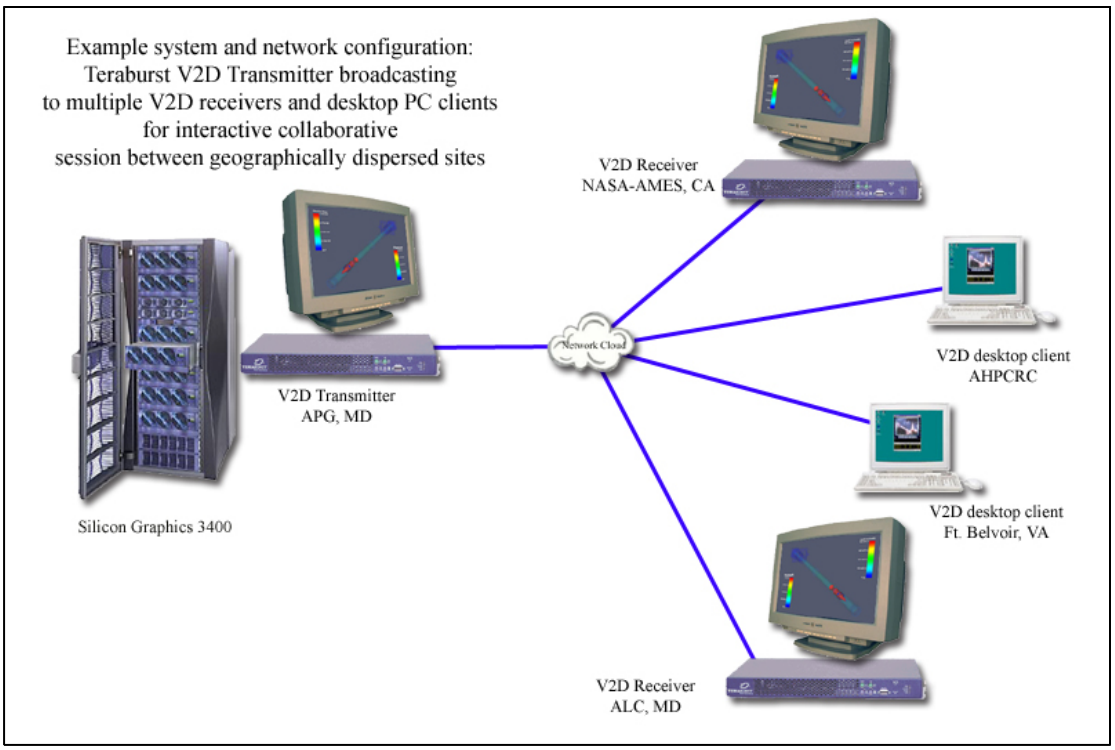

Figure 7. Example of the V2D system configuration with hardware and software clients.

*Teraburst Intelligent Frame Differencing is a trademark of Teraburst Networks, Inc.

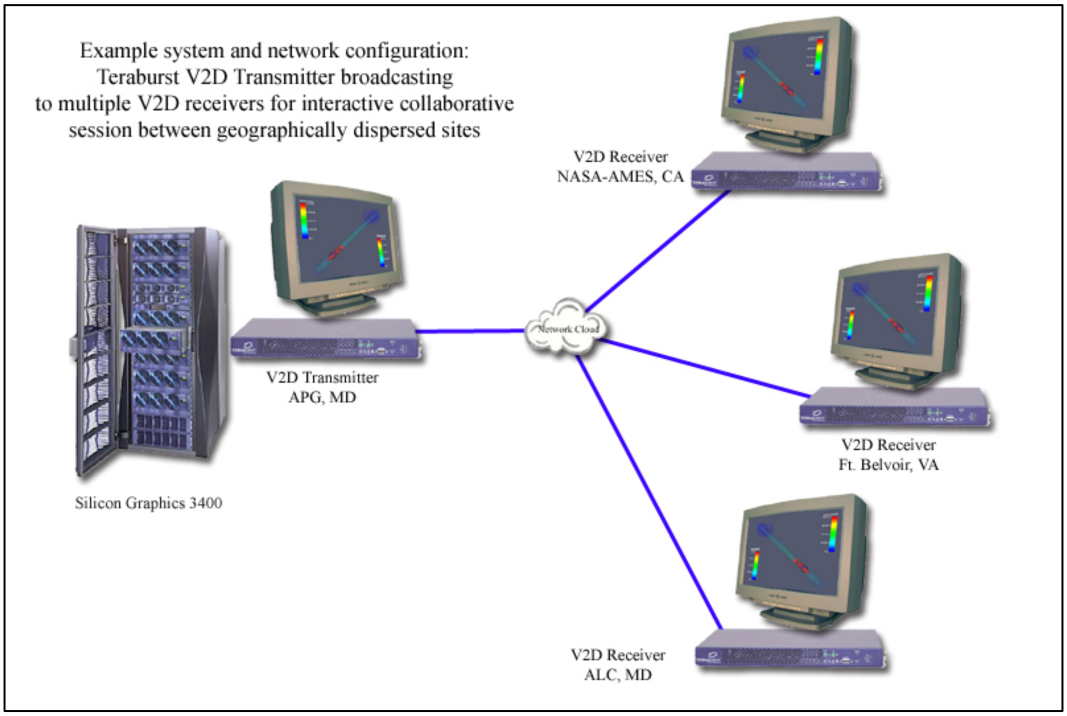

Figure 8. Example of a V2D system configuration using hardware clients.

These components can all introduce noise, leading to 0.5%–4% noise levels.[1] Each environment and installation will require some calibration to deal with inherent noise issues. Noise levels are also dependent upon types of frames; highly complex frames with areas of high contrast and high detail are more susceptible to noise than gradually changing images.[1]

By using Teraburst Intelligent Frame Differencing when static images are displayed, almost no information is sent across the network, and only new information is sent when moving images are displayed (figure 9).

Further reduction of bandwidth is achievable through spatial compression techniques.[1] The Teraburst Visualization Compression Engine[*] provides users with the flexibility to configure connections for desired and/or adequate levels of performance, allowing a user to select the trade-off between adequate picture quality and bandwidth utilization.

Compression ratios can be dialed in at 1:8–1:10 without visual loss. Higher compression rates up to 1:15 provide good results but introduce some artifacts. The Teraburst Visualization Compression Engine uses a combination of "4-4-4" and "4-2-2" subsampling. In 4-4-4 compression, there are equal amounts of luminance and chrominance information, while in 4-2-2 compression, half the chrominance information is used.[1]

[*]The Teraburst Visualization Compression Engine is a trademark of Teraburst Networks, Inc.

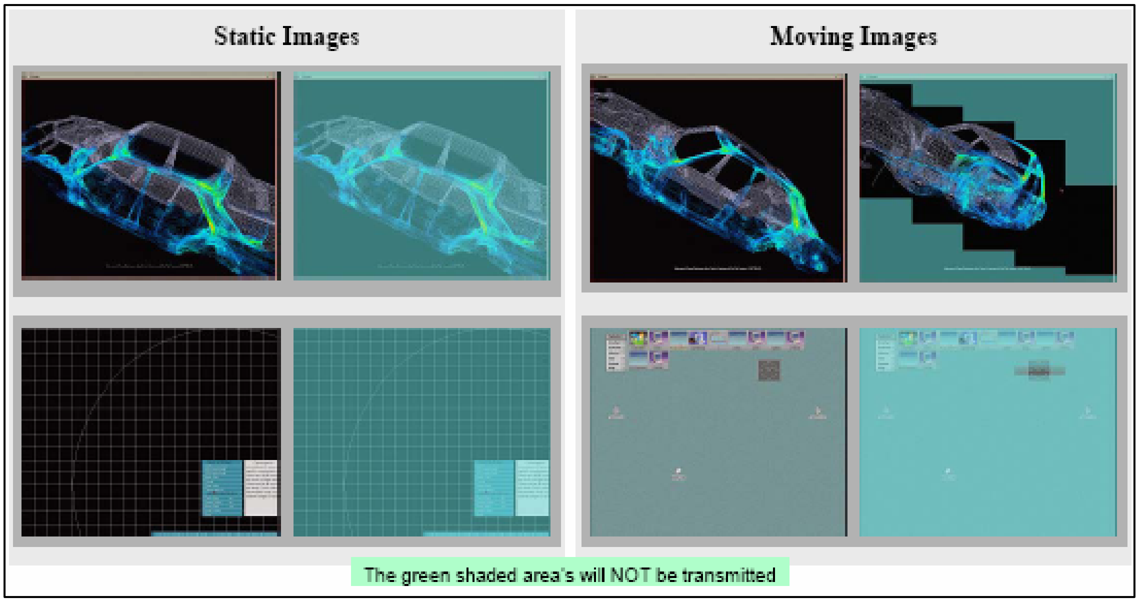

Source: Image provided by Teraburst Networks, Inc.

Figure 9. Static and moving images.

Teraburst has combined these compression methods for Ethernet connectivity of high-end visualization. The results are shown in table 1.

Table 1. Compression method results.

Resolution		**Compression Slice Drop**	**1:12 90%**	**1:8 75%**	**1:2 60%**	**Uncompressed No Drop**
Horiz	**Vert**	**Refresh Rate**	**(Mb/s)**	**(Mb/s)**	**(Mb/s)**	**(Mb/s)**
1280	1024	60	16 (fast-Ether)	59 (fast-Ether)	378 (gig-e)	1887
1280	1024	96	26 (fast-Ether)	95 (fast-Ether)	605 (gig-e)	3020
1600	1200	60	24 (fast-Ether)	87 (fast-Ether)	554 (gig-e)	2765
1600	1200	85	34 (fast-Ether)	123 (gig-e)	784 (gig-e)	3917
2048	1536	85	55 (fast-Ether)	202 (gig-e)	1285	6417

Note: Information provided by Teraburst Networks, Inc.

4. Future Direction

At present, Teraburst Networks, Inc. presents the most promising architecture for future expansion. The infrastructure of the V2D system supports the integration of several PCI video boards into one networked appliance for support of multiple graphic pipes on HPC assets. Teraburst also provides a PC thin client, which enables remote users to view V2D transmitter video streams. In the near future, keyboard and mouse interaction will be incorporated.

Visualization computing is also evolving due to the growing popularity of computing clusters. The Teraburst model provides the ability to integrate the V2D into the grid networking and is

proposing the use of “grid visualization” (figure 10). Besides providing remote access from everywhere, collaborating between any sites, and broadcasting to any place on the network, we can imagine a network of V2D receivers assembling different pieces of an image generated from different computer facilities on the grid.[1]

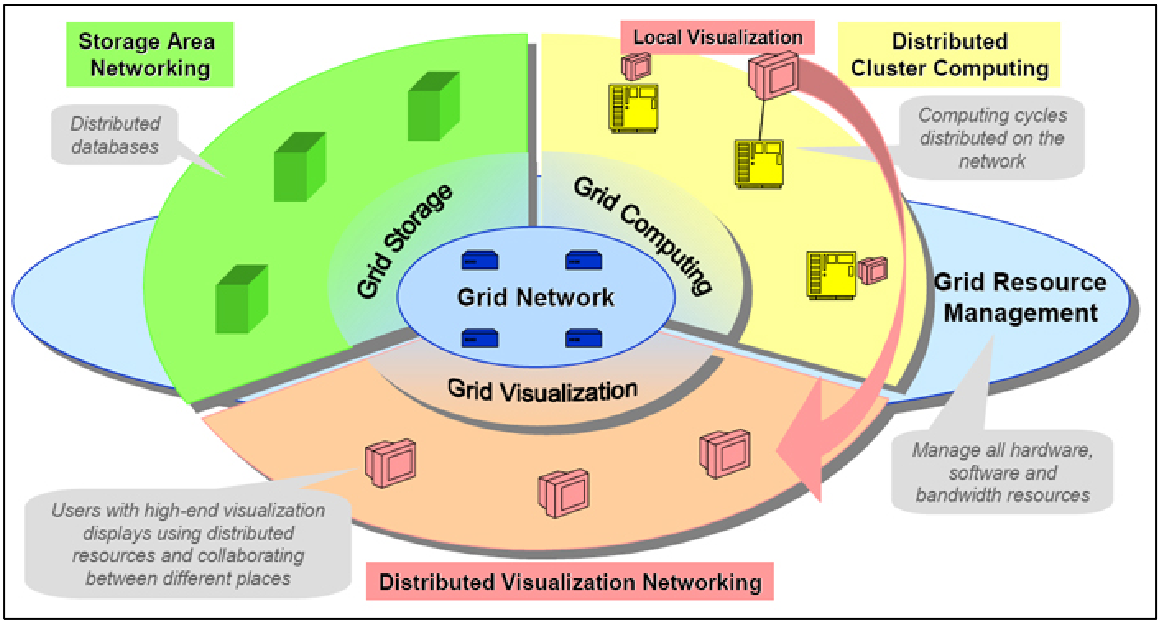

Source: Image provided by Teraburst Networks, Inc.

Figure 10. Distributed visualization networking.

NO. OF COPIES	ORGANIZATION
1 (PDF ONLY)	DEFENSE TECHNICAL INFORMATION CTR DTIC OCA 8725 JOHN J KINGMAN RD STE 0944 FORT BELVOIR VA 22060-6218
1	US ARMY RSRCH DEV & ENGRG CMD SYSTEMS OF SYSTEMS INTEGRATION AMSRD SS T 6000 6TH ST STE 100 FORT BELVOIR VA 22060-5608
1	INST FOR ADVNCD TCHNLGY THE UNIV OF TEXAS AT AUSTIN 3925 W BRAKER LN STE 400 AUSTIN TX 78759-5316
1	DIRECTOR US ARMY RESEARCH LAB IMNE ALC IMS 2800 POWDER MILL RD ADELPHI MD 20783-1197
3	DIRECTOR US ARMY RESEARCH LAB AMSRD ARL CI OK TL 2800 POWDER MILL RD ADELPHI MD 20783-1197
3	DIRECTOR US ARMY RESEARCH LAB AMSRD ARL CS IS T 2800 POWDER MILL RD ADELPHI MD 20783-1197

ABERDEEN PROVING GROUND

1	DIR USARL AMSRD ARL CI OK TP (BLDG 4600)

NO. OF COPIES	ORGANIZATION
	ABERDEEN PROVING GROUND
5	DIR USARL AMSRD ARL CI HC J VINES